Draft version July 8, 2016
Preprint typeset using LaTeX style AASTeX6 v. 1.0

arXiv:1602.06259v2 [astro-ph.IM] 6 Jul 2016

REDUNDANT ARRAY CONFIGURATIONS FOR 21 CM COSMOLOGY

Joshua S. Dillon[1,2,3] and Aaron R. Parsons[1,2]

[1]Department of Astronomy, UC Berkeley, Berkeley, CA
[2]Radio Astronomy Laboratory, UC Berkeley, Berkeley, CA
[3]Berkeley Center for Cosmological Physics, Berkeley, CA

ABSTRACT

Realizing the potential of 21 cm tomography to statistically probe the intergalactic medium before and during the Epoch of Reionization requires large telescopes and precise control of systematics. Next-generation telescopes are now being designed and built to meet these challenges, drawing lessons from first-generation experiments that showed the benefits of densely packed, highly redundant arrays—in which the same mode on the sky is sampled by many antenna pairs—for achieving high sensitivity, precise calibration, and robust foreground mitigation. In this work, we focus on the Hydrogen Epoch of Reionization Array (HERA) as an interferometer with a dense, redundant core designed following these lessons to be optimized for 21 cm cosmology. We show how modestly supplementing or modifying a compact design like HERA's can still deliver high sensitivity while enhancing strategies for calibration and foreground mitigation. In particular, we compare the imaging capability of several array configurations, both instantaneously (to address instrumental and ionospheric effects) and with rotation synthesis (for foreground removal). We also examine the effects that configuration has on calibratability using instantaneous redundancy. We find that improved imaging with sub-aperture sampling via "off-grid" antennas and increased angular resolution via far-flung "outrigger" antennas is possible with a redundantly calibratable array configuration.

1. INTRODUCTION

The quest to detect and characterize the 21 cm signal from neutral hydrogen during the Cosmic Dawn up through the Epoch of Reionization (EoR) has driven the design and construction of a generation of radio interferometers. The breadth of designs of these telescopes, which include the Low Frequency Array (LOFAR; van Haarlem et al. 2013), the Giant Metrewave Radio Telescope (GMRT; Paciga et al. 2013), the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), and the Murchison Widefield Array (MWA; Lonsdale et al. 2009; Tingay et al. 2013; Bowman et al. 2013), reflects different strategies for detecting the signal and separating it from bright astrophysical foregrounds. Because those foregrounds—mostly synchrotron emission from our Galaxy and comparatively nearby galaxies—are four to five orders of magnitude brighter than the cosmological signal (de Oliveira-Costa et al. 2008), their statistical separation is of paramount importance. Despite the difficulties, the scientific payoff from a successful observation will be considerable. The 21 cm signal contains vast amounts of information about formation and evolution of the first stars, galaxies, and black holes (Furlanetto et al. 2006; Morales & Wyithe 2010; Pritchard & Loeb 2012; Zaroubi 2013; Loeb & Furlanetto 2013) and could one day prove the most sensitive test of ΛCDM (McQuinn et al. 2006; Mao et al. 2008; Clesse et al. 2012).

As the next generation of 21 cm observatories like the Hydrogen Epoch of Reionization Array (HERA; Pober et al. 2014; DeBoer et al. 2016) and the low-frequency portion of the Square Kilometer Array (SKA-low; Koopmans et al. 2015) move from design to construction, it is essential to reflect upon the lessons of the first generation. The strategies they have employed to separate the signal from the foregrounds rely on the intrinsic spectral smoothness of the foregrounds. We know, however, that spectrally smooth foregrounds can create spectrally complicated three dimensional maps due to the intrinsic chromaticity of inteferometric measurements (Datta et al. 2010; Parsons et al. 2012b; Vedantham et al. 2012; Morales et al. 2012; Hazelton et al. 2013; Trott et al. 2012; Thyagarajan et al. 2013; Liu et al. 2014a,b). The effect, known as the "wedge" for the shape of foreground contamination in cylindrically binned power spectra (Pober et al. 2013), makes that separation harder. One can attempt to subtract a model

jsdillon@berkeley.edu

of foregrounds propagated through a model of one's instrument (Bernardi et al. 2011; Sullivan et al. 2012) and then subtract the remaining residuals statistically as the LOFAR team is attempting (Chapman et al. 2013; Bonaldi & Brown 2015), although this remains a considerable challenge.

One can also simply give up on modes within the maximal "horizon wedge" (corresponding to the delay of a point source at the horizon for a given baseline) and filter them out. This foreground avoidance strategy, championed by PAPER with its delay spectrum approach (Parsons et al. 2012b), sacrifices sensitivity (especially on the long baselines which exhibit the most spectral structure; Pober et al. 2014) for robust foreground excision and has enabled the most sensitive power spectrum limits in Parsons et al. (2014), Jacobs et al. (2015), and most recently in Ali et al. (2015). PAPER's highly redundant, compact configuration minimizes the cost of foregoing imaging by working only within the "EoR window," the region putatively uncontaminated by the foreground wedge.

The MWA, in contrast, uses its moderately dense configuration and imaging capability to pursue a hybrid approach (Dillon et al. 2014, 2015a; Trott et al. 2016), first subtracting foregrounds from maps (or visibilities) before projecting out the most contaminated wedge modes. In this case, the goal of foreground subtraction is not an increase in sensitivity by recovering more foreground contaminated modes, but rather to minimize the leakage of foreground contamination out of the wedge and into the EoR window (Dillon et al. 2015a; Pober et al. 2016). Since the foregrounds are so bright, any small leakage can swamp the 21 cm signal and lead to systematic errors that do not integrate down with additional observation. Subtracting foregrounds can both ameliorate those errors and get us closer to working within the wedge.

That said, a primary reason why spectrally smooth foregrounds can lead to contamination beyond the wedge is miscalibration (Shaw et al. 2015). Spectrally smooth foregrounds observed by an instrument with a subtly unsmooth spectral response can create foreground contamination inseparable from the cosmological signal. Traditionally, radio telescopes are calibrated by pointing at bright point sources with well-known spectra (Thompson et al. 2001). This is more complicated for instruments with the wide fields of view like LOFAR and especially the MWA, which must calibrate on a point source model with thousands of sources[1] and continually refine those models with new and deeper observations with better-understood maps (Braun 2013; Yatawatta et al. 2013).

In an array optimized for instantaneous redundancy like PAPER, the fact that each mode on the sky is measured multiple times allows for the simultaneous solution of visibilities and calibration parameters at a singe frequency (Wieringa 1992; Liu et al. 2010; Ram Marthi & Chengalur 2014). All that remains is bandpass calibration of the entire array, rather than each element individually. The redundant calibration approach, by not requiring a sky model for more than a single overall bandpass, is a powerful alternative to traditional radio inteferometric calibration. It was implemented successfully by Zheng et al. (2014) and proved a key improvement to PAPER analysis that enabled the Ali et al. (2015) limits.

The desire to maximize sensitivity on the short baselines that have the least wedge contamination by sampling them redundantly is generally thought to be at odds with high fidelity imaging. It seems, on the face of it, that we must choose between the two strategies for calibration and foreground mitigation. This paper seeks to answer the question: can we have our cake and eat it too? Can we build a redundant array with high sensitivity on short baselines that enables both strategies for calibration and foreground mitigation? To answer that question in a specific case, we start with the core design for HERA—331 hexagonally packed 14 m dishes (Neben et al. 2016b; Ewall-Wice et al. 2016c; Thyagarajan et al. 2016; Patra et al. *in prep.*)—and examine how modest additions or rearrangements of antennas can improve mapmaking ability without sacrificing the precise calibratability or high sensitivity enabled by redundancy.

Previous array configuration studies have focused on other goals. Pen et al. (2004), Lonsdale et al. (2009), and Wijnholds et al. (2011) look at array configurations designed to suppress sidelobes, but do not consider redundancy or optimizing EoR window sensitivity. Tegmark & Zaldarriaga (2010) explore the broad class of array designs that can be correlated with the fast Fourier transform (FFT) but did not pursue any particular optimization scheme. Parsons et al. (2012a) explore the benefits of instantaneous redundancy, but do not optimize for redundant calibratability or imaging capability. More recently, Greig et al. (2015) focus on varying levels of array compactness and their effects on sensitivity as manifested by errors on a parametrized model of reionization, but does not consider the benefits of redundancy for controlling systematics.

In this work, we assess the mapmaking capability and redundant calibratability of three variants of the HERA configuration, developing criteria for assessing array configurations that address not just sensitivity but

[1] Observed with a beam power pattern that may vary from time to time and antenna to antenna, necessitating simultaneous beam and antenna gain calibration.

also the all-important control of systematic effects. In Section 2, we discuss the motivations for the design of the HERA core and present variants that increase its ability to make widefield, high-resolution maps without sacrificing sensitivity. In Section 3, we look in detail at both the qualitative and quantitative effects that array configuration has on mapmaking. Next, in Section 4, we examine the relative redundant calibratability of the arrays, including the distant "outrigger" antennas. Finally, in Section 5 we conclude with a discussion of the lessons learned for HERA in particular and for the next generation of 21 cm interferometers more generally.

2. ARRAY DESIGN

We begin our analysis of modified HERA configurations by presenting three variants to the HERA core design, a hexagonally packed core of 14 m dishes (see Figure 1 and, for more detail, DeBoer et al. 2016). The need to achieve a large collecting area cheaply while maximizing the sensitivity to short baselines drives this basic design, since the benefits of coherent averaging generally outweigh the additional sample variance due to a limited field of view (Parsons et al. 2012a; Pober et al. 2014). In fact, its 21 cm focused design allows it a similar performance to the SKA-low, especially when foregrounds must be avoided rather than subtracted (Pober et al. 2014; Greig & Mesinger 2015; Ewall-Wice et al. 2016a). HERA's zenith-pointing dishes, while restricting the survey area to a single stripe in declination, minimize the so-called "pitchfork effect" (Thyagarajan et al. 2015c,b) in which diffuse emission near the horizon contributes disproportionately to the power that leaks into the EoR window.

The array configuration has a number of advantages that we would like to preserve:

1. Its compactness enables it to achieve a high - sensitivity detection of the 21 cm signal using only foreground avoidance.
2. It can be calibrated precisely using its redundancy.
3. Its hexagonal configuration makes FFT correlation possible (Tegmark & Zaldarriaga 2009, 2010), although it is not strictly necessary for an array with only a few hundred elements.

It also has a number of drawbacks:

1. It has fairly low angular resolution.
2. It measures only a relatively small number of modes simultaneously—eventually, sample variance dominates these measurements.
3. The modes it does measure fall on a regular hexagonal grid, leading to grating lobes in the point spread function (PSF).

In this section, we identify strategies for mitigating these shortcomings with minor modifications to the array configuration. In particular, we asses methods of subsampling the hexagonal baseline grid (Section 2.1) and for increasing the angular resolution of the array by adding far-flung "outrigger" antennas (Section 2.2). Then, we show that all these designs have approximately the same raw power spectrum sensitivity.

2.1. *Subsampling the Baseline Grid*

Although a 331-element hexagonal core measures 54,615 visibilities at any one time, it actually only measures 630 unique separations between antennas. These fall on a hexagonal grid, since the convolution of a hexagon with itself is simply a larger hexagon. In a compact array, the minimum separation between antennas, 14.6 m, sets the minimum baseline length and the fundamental size of that grid of measured baselines. And since the Fourier sky is regularly sampled at spacings much wider than Nyquist sampling ($\lambda/2$), we observe grating lobes in the PSF due to signal aliasing caused by undersampling. If we want to suppress grating lobes, we need denser sampling and more independent information.

Of course, our visibility measurements are not disconnected δ-functions in Fourier space. The visibility, which is a measurement of the correlation of voltages measured by two antennas separated by baseline vector $\mathbf{b}$, is given by

$$V(\mathbf{b},\nu) = \int B(\hat{\mathbf{s}},\nu)\, I(\hat{\mathbf{s}},\nu) \exp\left[-2\pi i \frac{\nu}{c}\mathbf{b}\cdot\hat{\mathbf{s}}\right] d\Omega. \quad (1)$$

where $I(\hat{\mathbf{s}},\nu)$ is the sky intensity as a function of position and frequency and $B(\hat{\mathbf{s}},\nu)$ is the product of the direction and frequency dependence of the response of individual antennas.[2]

For a coplanar array sampling the sky instantaneously, we can think of a visibility as a average over a region of the Fourier transformed sky, projected onto a plane and convolved with the Fourier transform of the beam product (Thompson et al. 2001). Alternatively, we can think of it as a δ-function sampling of the beam-convolved sky. HERA's beams are roughly axisymmetric and therefore our weighted average is over an approximately circular region of Fourier space (also known as the uv-plane). It follows that we have the least information about the parts the uv-plane furthest from the exact baselines probed. If we would like to make new measurements inside the core (in uv-space) as independent as possible (although not, of course, completely independent of existing baselines), then we would like to place antennas so

[2] For which, in this work, we use a simulation of actual HERA elements (Neben et al. 2016b; Ewall-Wice et al. 2016c; DeBoer et al. 2016).

as to produce new baselines at the centroids of the lattice of equilateral triangles that make up the hexagonal grid.

For the sake of comparison, we start with a "fiducial" HERA core, which we call configuration (a) and show in the first panel of Figure 1. In Figure 2 we can see that hexagonal grid pattern of instantaneously redundant baselines (as opposed to baselines that measure the same mode at a later time due to the rotation of the earth.) We note that the best-sampled baselines are the shortest. Longer separations have fewer corresponding pairs of antennas within the core.

To produce new baselines at the centroids of the triangular lattice, we need to add antennas that are not on the hexagonal grid or dither existing antennas. If we pick any two 14.6 m basis vectors that define the hexagonal grid in linear, integer combinations (e.g. up-and-to-the-left and up-and-to-the-right), we note the centroids fall at 1/3 or 2/3 of the sum of those two vectors. In other words, by introducing antennas dithered from the main grid by 1/3 or 2/3 of that separation, we can add a grid of baselines that effectively subsample the main hexagonal grid. Maximally packed dishes Nyquist sample modes within the main lobe of the primary beam, but they produce grating lobes in the sidelobes (which would require $\lambda/2$ spacing to eliminate for the whole sky). Sub-aperture sampling suppresses these grating lobes, as we will see in Section 3.

We employ this strategy in developing our two alternative array configurations, (b) and (c), shown in Figure 1. In (b), we add a small number of antennas near the perimeter of the main hexagon, each displaced from the main grid. In (c) we split the entire hexagon into three sectors, each relatively displaced by those same increments. Both patterns allow for complete sampling at triple the density of the original grid (see Figure 2). In the case of the split-core configuration (c), that triple-density coverage does not extend to the full hexagonal core, but those baselines can be filled in by outrigger antennas, as we will discuss in Section 2.2.[3]

An important contrast between configurations (b) and (c) is that while off-grid baselines in (b) have at most a handful of measurements, the distinction between "on-grid" and "off-grid" in (c) becomes less clear for longer baselines. That is because longer baselines are more likely to be inter-sector and thus off-grid while shorter baselines are more often intra-sector and thus on-grid. The result is more distributed and thus uniform cover-

[3] As an aside, the bottom row of (c) is not necessary for the desired *uv*-coverage and can be eliminated as a cost saving measure. In this paper, we keep it to make the cores more directly comparable.

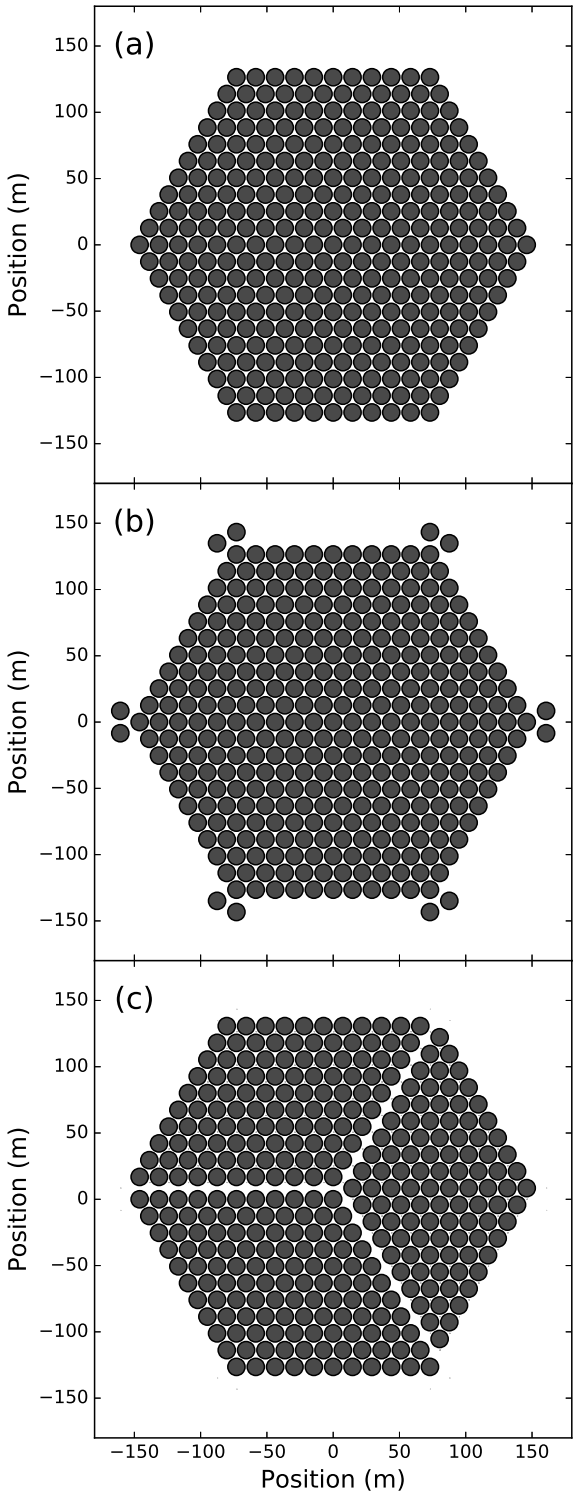

Figure 1. Three different versions of the HERA core, meant to produce different instantaneous baseline distributions. In (a) we take the standard hexagon, which produces a hexagonal grid of baseline separations (see Figure 2). In (b) we add pairs of antennas that are displaced 1/3 or 2/3 of the sum of two hexagonal basis vectors (e.g. right and up-and-to-the-right). In (c) we fracture the entire array into three sectors, displaced 0/3, 1/3, or 2/3 of the sum of two hexagonal basis vectors. The configurations have 331, 343, and 331 antennas, respectively. These different options affect both mapmaking and calibratability, as we discuss in Sections 3 and 4.

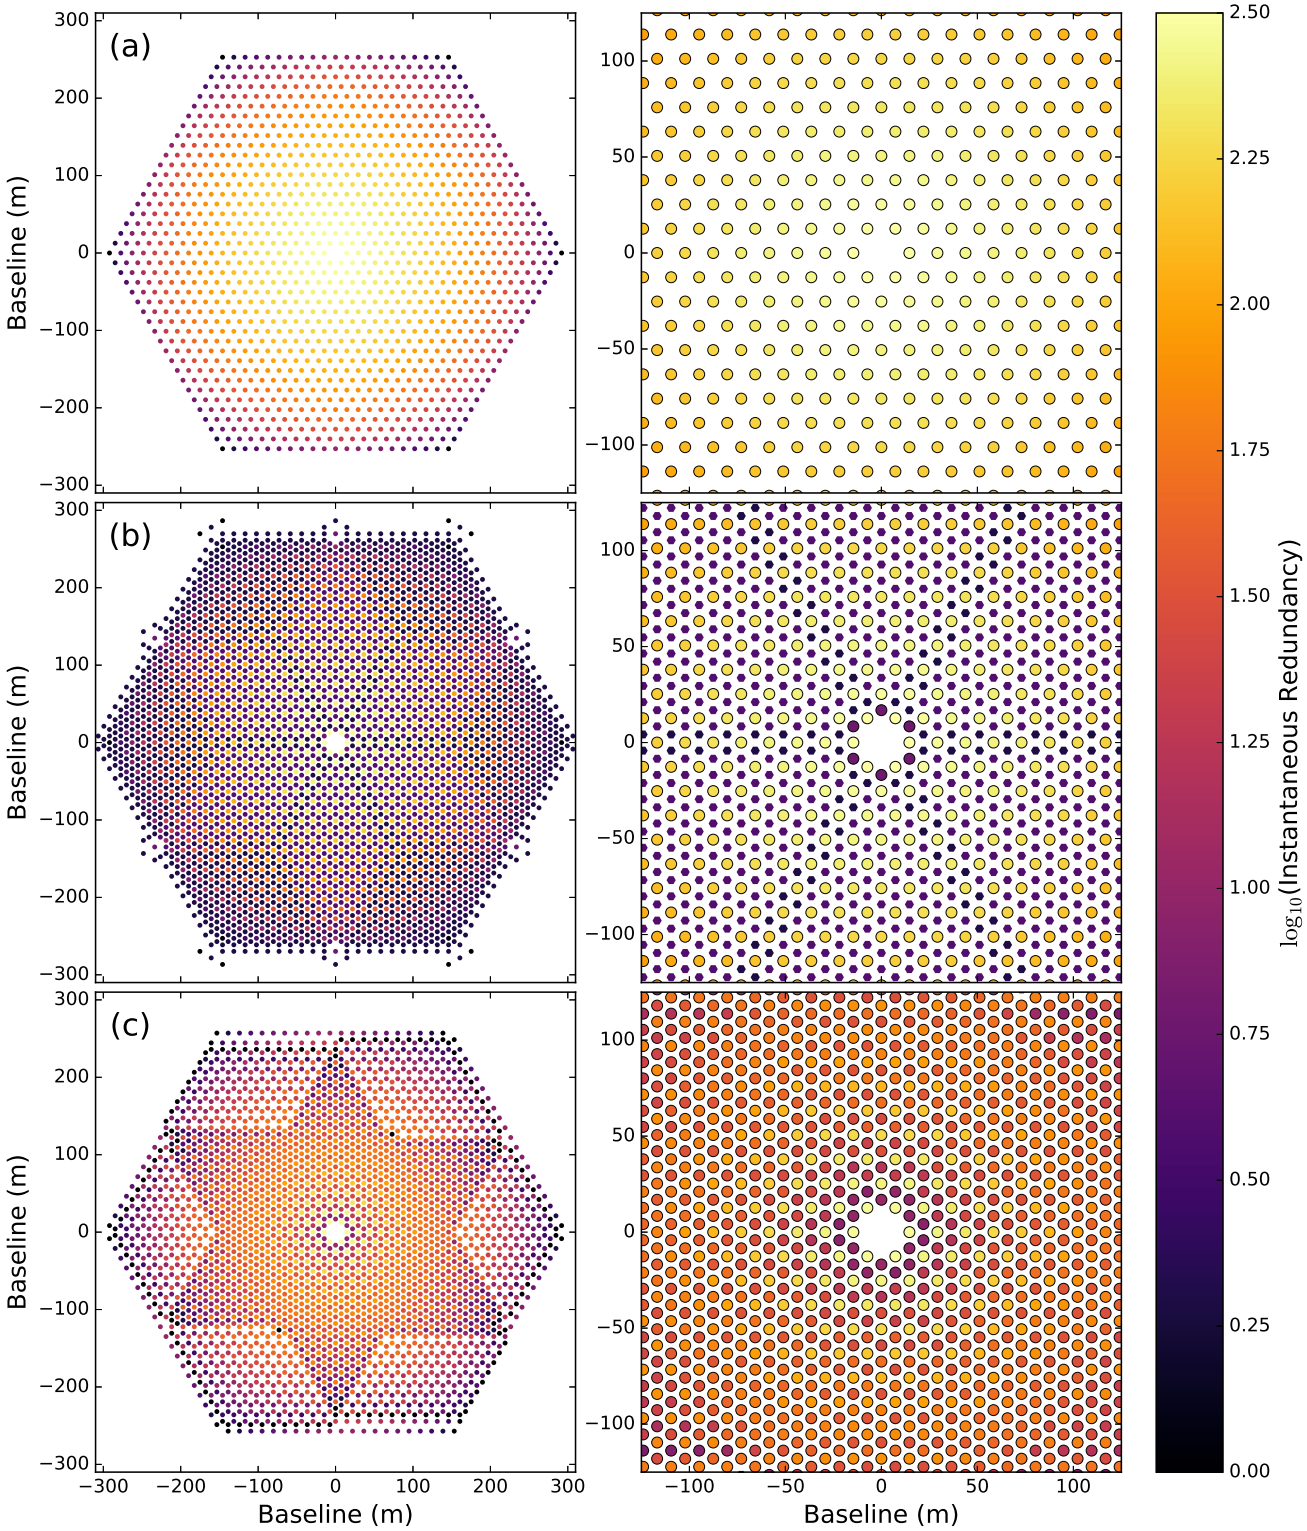

Figure 2. The baseline distributions produced instantaneously by the three different array configurations in Figure 1. Color indicates how many antenna pairs measure the given separation vector at any one time. On the left we show the full baseline pattern, on the right we zoom into the highest sensitivity region. In the right panels, baselines sampled four or fewer times are indicated by smaller markers in order to better illustrate the redundancy pattern. An unmodified core in (a) produces a hexagonal grid of well-sampled baselines. Configuration (b) adds off-grid antennas to each corner of the core to produce baselines that effectively subsample the hexagonal grid. The splitting of the core in (c) achieves the same sampling pattern as in (b) without any additional antennas. In contrast to (b), the off-grid baselines in (c) are more frequently and evenly sampled. That is because longer baselines are usually between relatively displaced sectors. This also leads to unsampled points on the grid outside the main hexagram, but that deficiency can be rectified with outrigger antennas (see Figure 4). The size of the points is not physical; individual baselines measure δ-function samples of the beam-convolved sky. Therefore, the sub-aperture samples of configurations (b) and (c) are not measuring completely independent sky information from their neighboring baselines.

age of the triple-density grid than in (b).[4] It should be noted that configurations (a) and (b), by virtue of their denser cores, concentrate somewhat more sensitivity on shorter baselines, which are those less affected by the wedge. We also note that (b) and (c) do not substantially affect the possibility of performing FFT correlation, since it merely requires that all antennas fall on a regular or hierarchically regular grid (Tegmark & Zaldarriaga 2010); gridding based FFT schemes (Morales 2011; Thyagarajan et al. 2015a) are not necessary. That grid is now simply finer and has many unobserved vertices that must be flagged.

2.2. *Increasing Angular Resolution*

In addition to the hexagonal cores, we are interested in supplementing each design with several far-flung "outrigger" antennas in order to improve angular resolution (see Figure 3). Adding a single antenna far from the core adds an entire hexagon of baselines to the *uv*-plane. This improves the imaging capability of the array and can help with foreground removal (see Section 3). With configuration (a), we pick the minimal set of additional antennas to tile the maximum area of the *uv*-plane with hexagons. Although hexagons do tessellate, the borders between them create discontinuities in the hexagonal grid, as we can see in Figure 4. Tessellating hexagons produces some overlapping regions in the core, but outside the core the scheme is extensible to longer baselines with no additional overlap.

Configuration (b) takes inspiration from configuration (a) by also using the hexagonal perimeter of the core to tile the *uv*-plane. Its outrigger antennas differ in two notable ways. First, it includes outrigger antennas on both sides of the core. Second, that grid is contracted slightly from the grid in (a) such that there exist redundant baselines measured between different outrigger antennas. Both these changes are designed to allow for redundant calibratability, which we will discuss in Section 4.

Finally, configuration (c) takes a different approach. While it too is designed for redundant calibratability, it takes advantage of the split-core configuration to create full baseline coverage at triple density (see Figure 4). Compared to (a) or (b), that increased density allows it to better suppress the sidelobes of point sources at the cost of some angular resolution.

2.3. *Raw Power Spectrum Sensitivity*

[4] We considered carrying this further and splitting the hexagon into nine sectors to create even finer gridding. It turns out that it is difficult to make such a configuration with full *uv*-coverage and they appear to be less redundantly calibratable (and probably more difficult to build). We therefore do not consider such a design in this work.

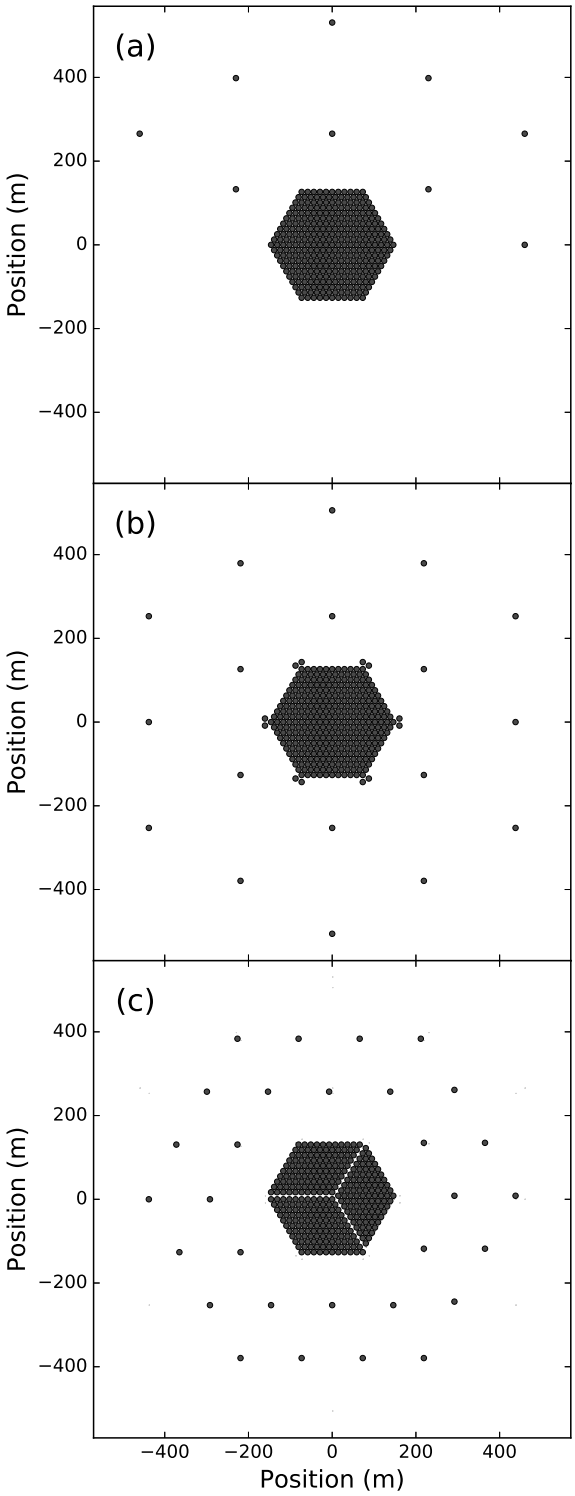

Figure 3. Three different approaches to "outrigger" antennas that increase the angular resolution of the array core designs in Figure 1. The outrigger antennas in (a) and (b) are designed to tessellate hexagonal sampling regions formed by outrigger-core antenna pairs in order to efficiently maximize the region of full coverage and thus increase angular resolution. However, (b) includes antennas on both sides of the core and is slightly more compact than (a), enhancing redundant calibratability (see Section 4). By contrast, the outrigger antennas in (c) are designed for both redundant calibration and to extend triple-density *uv*-coverage to longer baselines (see Figure 4). They have 340, 361, and 361 antennas, respectively.

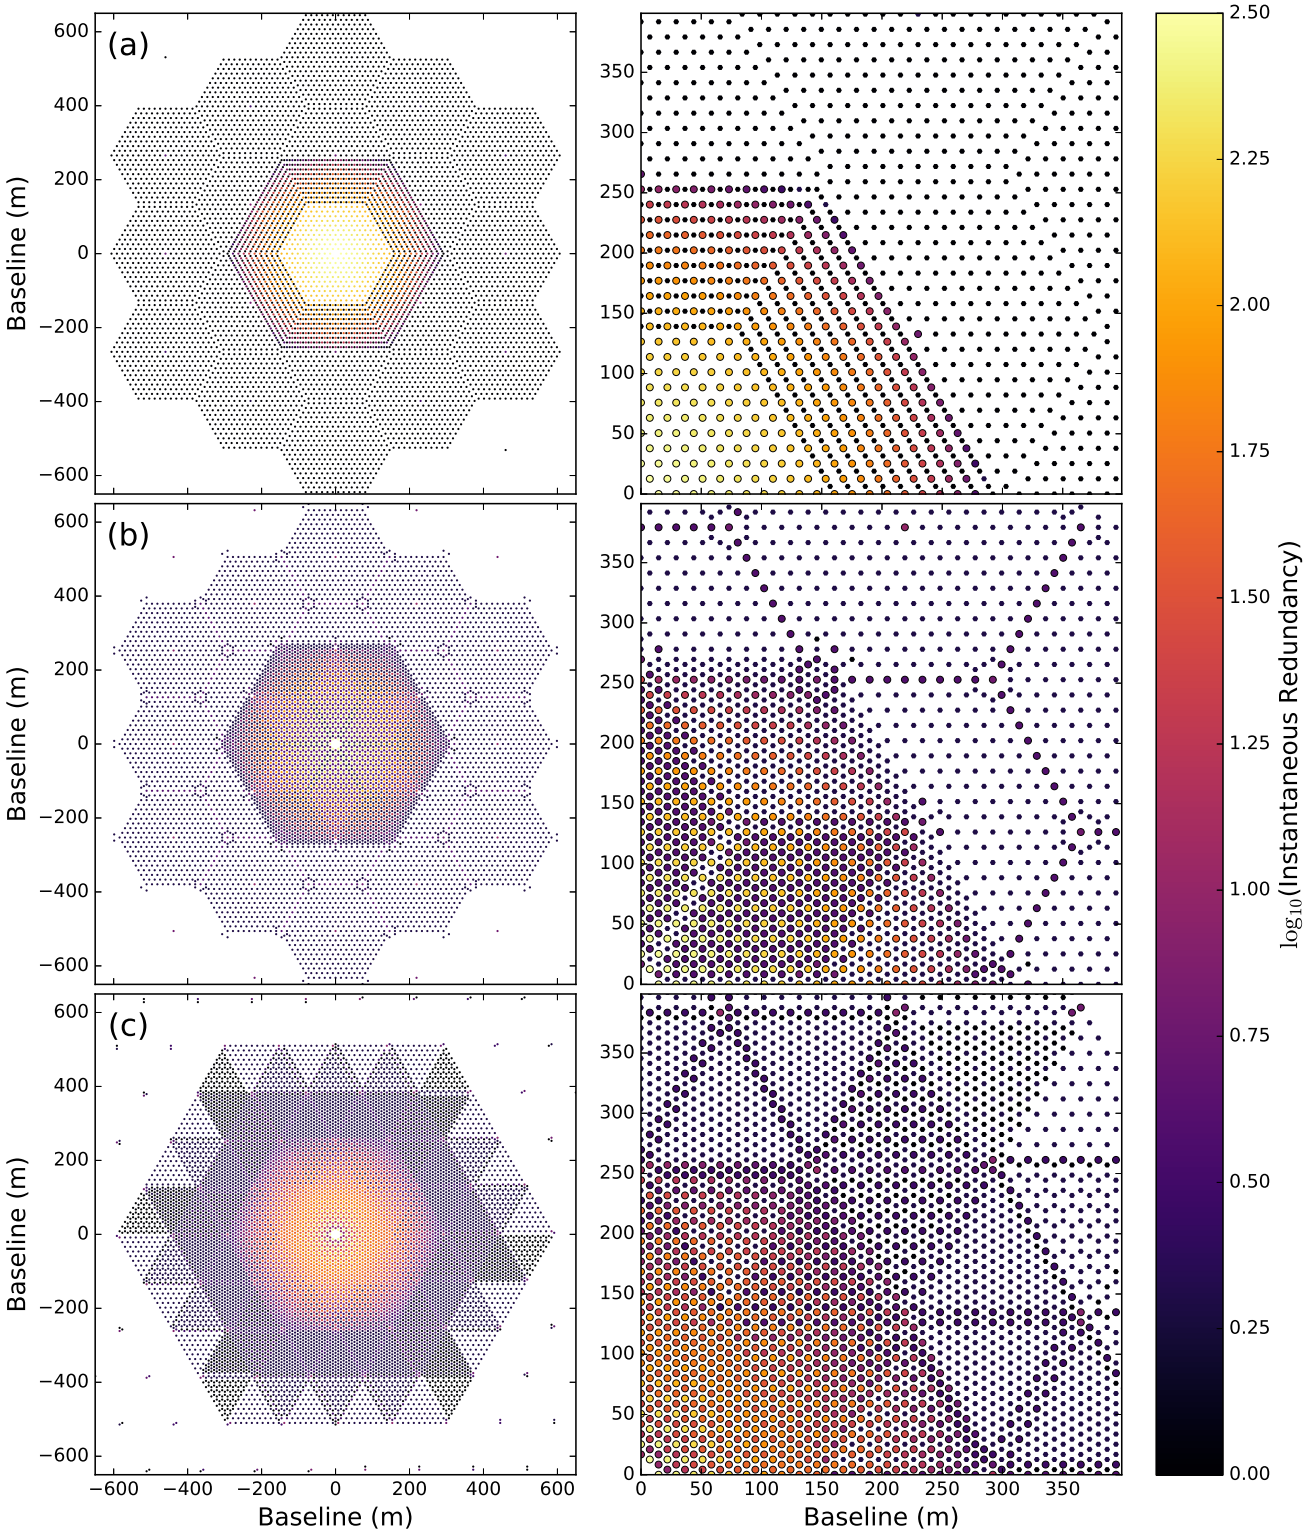

Figure 4. Instantaneous baseline sampling produced by the three different outrigger configurations in Figure 3. On the left we show the full baseline sampling pattern. On the right we show a zoom into a representative section. Here baselines sampled two or fewer times are denoted by smaller markers in the righthand panels in order to better illustrate the network of redundancy. In particular, we can see how contracting the outrigger positions from configuration (a) to configuration (b) produces overlapping edges of redundantly sampled baselines, creating a network by which the whole array can be redundantly calibrated (see Section 4). By contrast, the design in (c) arranges baselines to achieve complete coverage at triple density to as large a radius as possible while also maintaining plenty of overlap for redundant calibration. Additional outrigger antennas can be added to produce expanded *uv*-coverage for higher angular resolution imaging.

Before we move on to compare these configurations in terms of mapmaking and calibration, it is useful to check that these designs do not differ significantly in terms of their raw sensitivity to the 21 cm power spectrum. Since most sensitivity comes from short baselines and all three configurations have very similar cores, we do not expect much variation.

A proper calculation of the significance with which a fiducial model of reionization can be detected by these arrays is quite difficult. In addition to depending on foreground removal strategies (Pober et al. 2014) and models for the covariance of foreground residuals (Dillon et al. 2015a), one should take into account the partial coherence of baselines that measure overlapping modes in the uv plane, either from the same baseline at different times or from sub-aperture sampling. This is possible, though computationally challenging. No algorithm exists yet that propagates these effects into the power spectrum (Dillon et al. 2015b). A full treatment is therefore outside the scope of this paper.

Instead, we follow Pober (2015) and employ the `21cmSense` code[5] to explore the sensitivities of three configurations. This is quantified in terms of a "cumulative detection significance," which is the square root of the sum in quadrature of the ratio of the 21 cm power spectrum to the thermal noise over all wavenumbers k. In particular, we look at an 8 MHz bandwidth centered on 135 MHz for a 1080 hour drift scan and calculate the sensitivity to a model power spectrum from Mesinger et al. (2011) with a midpoint of reionization at $z = 9.5$. Table 1 shows that, regardless of our foreground mitigation strategy, all three of our arrays perform very similarly. We therefore should consider other criteria to assess their relative merit.

3. MAPMAKING CAPABILITY

The modifications and additions to the core HERA configuration that we presented in Section 2 were largely motivated by the desire to increase the density and extent of baseline sampling of the uv-plane in order to supplement HERA's imaging capability. It is worthwhile therefore to assess the ways in which imaging capabilities are different between the different configurations. As we alluded to earlier, we expect the outrigger antennas and off-grid antennas to affect angular resolution and widefield imaging, respectively. We can assess these effects qualitatively by comparing PSFs, which we do in Section 3.2. We can also be more quantitative in our comparison by looking at the expected noise and noise correlations in the final maps, as we do in Section 3.3.

[5] First developed for Pober et al. (2014) and now available publicly at `https://www.github.com/jpober/21cmSense`.

Table 1. Estimated detection significance for the various configurations.

Configuration	$N_{\rm ant}$	$N_{\rm core}$	Mod. S/N	Opt. S/N
Fiducial (a)	340	331	24.4σ	91.9σ
Corners (b)	361	343	25.0σ	95.2σ
Split-core (c)	361	331	24.0σ	93.7σ

NOTE—This calculation is performed for a reionization model that reaches 50% ionization at $z = 9.5$ (Mesinger et al. 2011) for each array configuration. We examine both a moderate scenario where foregrounds must be avoided outside the horizon window (plus a $0.1\,h\,{\rm Mpc}^{-1}$buffer), and an optimistic scenario where foregrounds outside the primary beam wedge can be subtracted. This approximate treatment using `21cmSense` does not account for partial coherence between baselines that sample overlapping regions of the uv plane, and therefore affords no advantage for the off-grid baseline sampling of configurations (b) and (c). Instead the results are dominated by the number antennas, $N_{\rm ant}$, and (especially for the moderate foreground scenario) by their relative compactness (which is related to $N_{\rm core}$). These relatively small differences in sensitivity are unlikely to impact suitability for 21 cm cosmology.

But first, we begin in Section 3.1 with a review of the optimal mapmaking formalism of Dillon et al. (2015b) that underlies the rest of the detailed mapmaking comparisons in this work.

3.1. *Mapmaking Review*

Fundamentally the problem of mapmaking is one of data reduction. We have a large quantity of time-ordered measurements, expressed as a data vector $\mathbf{y}$, which sample different linear combinations of the sky at different times and in different ways. However, those data are noisy, so we must be cognizant of how we weight and combine them.

The action of the interferometer measuring visibilities (as in Equation 1) can be expressed as a relationship between some discretized sky $\mathbf{x}$ and our measurements $\mathbf{y}$ as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}. \tag{2}$$

Here $\mathbf{A}$ represents the interferometric measurement—the family of integrals like Equation 1 but discretized—and $\mathbf{n}$ is the thermal noise. The discretization of $\mathbf{A}$ is demonstrated in Dillon et al. (2015b), which also includes a more detailed discussion of the following definitions and derivations.

If the statistics of the thermal noise are Gaussian and we define $\mathbf{N} \equiv \langle \mathbf{n}\mathbf{n}^\dagger \rangle$, then the optimal estimator (in the sense that no information about the sky is lost), $\widehat{\mathbf{x}}$, is

$$\widehat{\mathbf{x}} = \mathbf{D}\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{y} \tag{3}$$

where $\mathbf{D}$ is some invertible normalization matrix (Tegmark 1997). To understand the statistics of this estimator, we should look at its mean and covariance.

Since the sky is not random and since the thermal noise has $\langle \mathbf{n} \rangle = 0$,

$$\begin{aligned}\langle \hat{\mathbf{x}} \rangle &= \langle \mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}(\mathbf{A}\mathbf{x} + \mathbf{n}) \rangle \\ &= \mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}(\mathbf{A}\mathbf{x} + \langle \mathbf{n} \rangle) \\ &= \mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}\mathbf{x} = \mathbf{P}\mathbf{x}, \end{aligned} \tag{4}$$

where we have defined

$$\mathbf{P} \equiv \mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}. \tag{5}$$

The matrix $\mathbf{P}$ is the matrix of PSFs that tells us how a source in any pixel in $\mathbf{x}$ will appear in all other pixels. Likewise the covariance of the estimator is given by

$$\begin{aligned}\mathbf{C} &= \left\langle (\hat{\mathbf{x}} - \langle \hat{\mathbf{x}} \rangle)(\hat{\mathbf{x}} - \langle \hat{\mathbf{x}} \rangle)^\dagger \right\rangle \\ &= \left\langle (\mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{n})(\mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{n})^\dagger \right\rangle \\ &= \mathbf{D}\mathbf{A}^\dagger \mathbf{N}^{-1} \left\langle \mathbf{n}\mathbf{n}^\dagger \right\rangle \mathbf{N}^{-1}\mathbf{A}\mathbf{D}^\dagger \\ &= \mathbf{P}\mathbf{D}^\dagger. \end{aligned} \tag{6}$$

Since the only random thing here is the noise, this $\mathbf{C}$ is really the noise covariance of every map pixel with every other map pixel.

For estimating power spectra, Dillon et al. (2015b) argued that choosing a simple diagonal form of $\mathbf{D}$ is ideal for propagating mapmaking statistics into the power spectrum. In this work, we simply take $\mathbf{D} \propto \mathbf{I}$ such that the PSF of a source at zenith peaks at 1. In the flat-sky approximation used by most radio interferometers with small fields of view, this choice is analogous to natural weighting, where every independent measurement is weighted only by the noise in that measurement and not by the number of other measurements that fall in the same uv-cell.

However, it is also useful to consider the other extreme choice for $\mathbf{D}$. If we wanted an unbiased estimator of the sky, one where $\langle \hat{\mathbf{x}} \rangle = \mathbf{x}$, then we should choose $\mathbf{D} = (\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A})^{-1}$ in order to obtain $\mathbf{P} = \mathbf{I}$. In the flat-sky approximation, this choice is analogous to uniform weighting of cells in the uv-plane. With a limited field of view (low uv resolution) and limited angular resolution (small maximum u and v), it is possible to sample every uv-cell and produce a δ-function PSF using uniform weighting. This generally produces very noisy maps. And of course, depending on the pixelization, $\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}$ is often not invertible and this procedure becomes impossible without further assumptions about the sky. However, when this procedure is possible, it follows then that

$$\mathbf{C} = \mathbf{P}\mathbf{D}^\dagger \longrightarrow (\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A})^{-1}. \tag{7}$$

Since inverse covariance weighting for a future data reduction step involves multiplication by $\mathbf{C}^{-1} = \mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}$, this means that $\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}$ measures the information content in our maps. In Section 3.3, we will explore the structure of this matrix for our array configurations.

3.2. *PSF Comparison*

We start with a more qualitative comparison of the PSFs of the three configurations. In Figure 5, we examine the central region of a high-resolution HEALPix (Górski et al. 2005) map of the PSF of a source in the center of a 20° facet ($N_{\rm side} = 256$). At 150 MHz, that 20° corresponds to roughly double the size of HERA's primary beam. We examine the PSF produced by a single 10 second snapshot with the facet center at zenith and by a 40 minute drift scan observation centered on the facet. We also examine two weighting schemes, one where $\mathbf{N}$ simply reflects noise on a given redundantly sampled baseline (e.g. "natural" weighting) and another where no redundant baseline is weighted more than 50 times higher than any other. This is akin to "robust" weighting (Thompson et al. 2001), which transitions from natural weighting of low signal-to-noise ratio (S/N) baselines to uniform weighting of high S/N baselines, balancing noise against sidelobe confusion.

At first glance, the PSFs are very similar. This is to be expected, especially for the naturally weighted PSFs, which are so dominated by the core baselines. The prominent six-fold sidelobes we see in all panels of Figure 5 are due to the inter-element spacing and configuration in the core.

There are some small differences between the configurations. The main lobe of the PSF is slightly more sharply defined in configurations (a) and (b) because their uv-coverage extends to longer baselines (see Figure 4). More prominently, the sidelobes are dimmer and more smeared out in configuration (c) than in either (a) or (b). This makes sense since we expected the sidelobes to be suppressed in the split-core configuration, which gives more weight to the off-grid baselines. Our "robust" weighting scheme helps us see this more clearly, since it gives relatively more weight to off-grid baselines and better demonstrates the sidelobe suppression possible with configuration (c).

3.3. *Quantitative Mapmaking Results*

The observed sidelobe suppression in Figure 5 raises a question: does configuration (c) actually have an improved ability to measure modes far from the center of the primary beam "more independently"? To answer that, we need to analyze the structure of the matrix $\mathbf{A}^\dagger \mathbf{N}^{-1}\mathbf{A}$, the noise-weighted instrument response which contains the information and correlations in our map. Were it computationally feasible, we would do this for a full-sky, high- resolution map. However, the numerical difficulty of applying the rigorous mapmaking treatment of Section 3.1 necessitates some approxima-

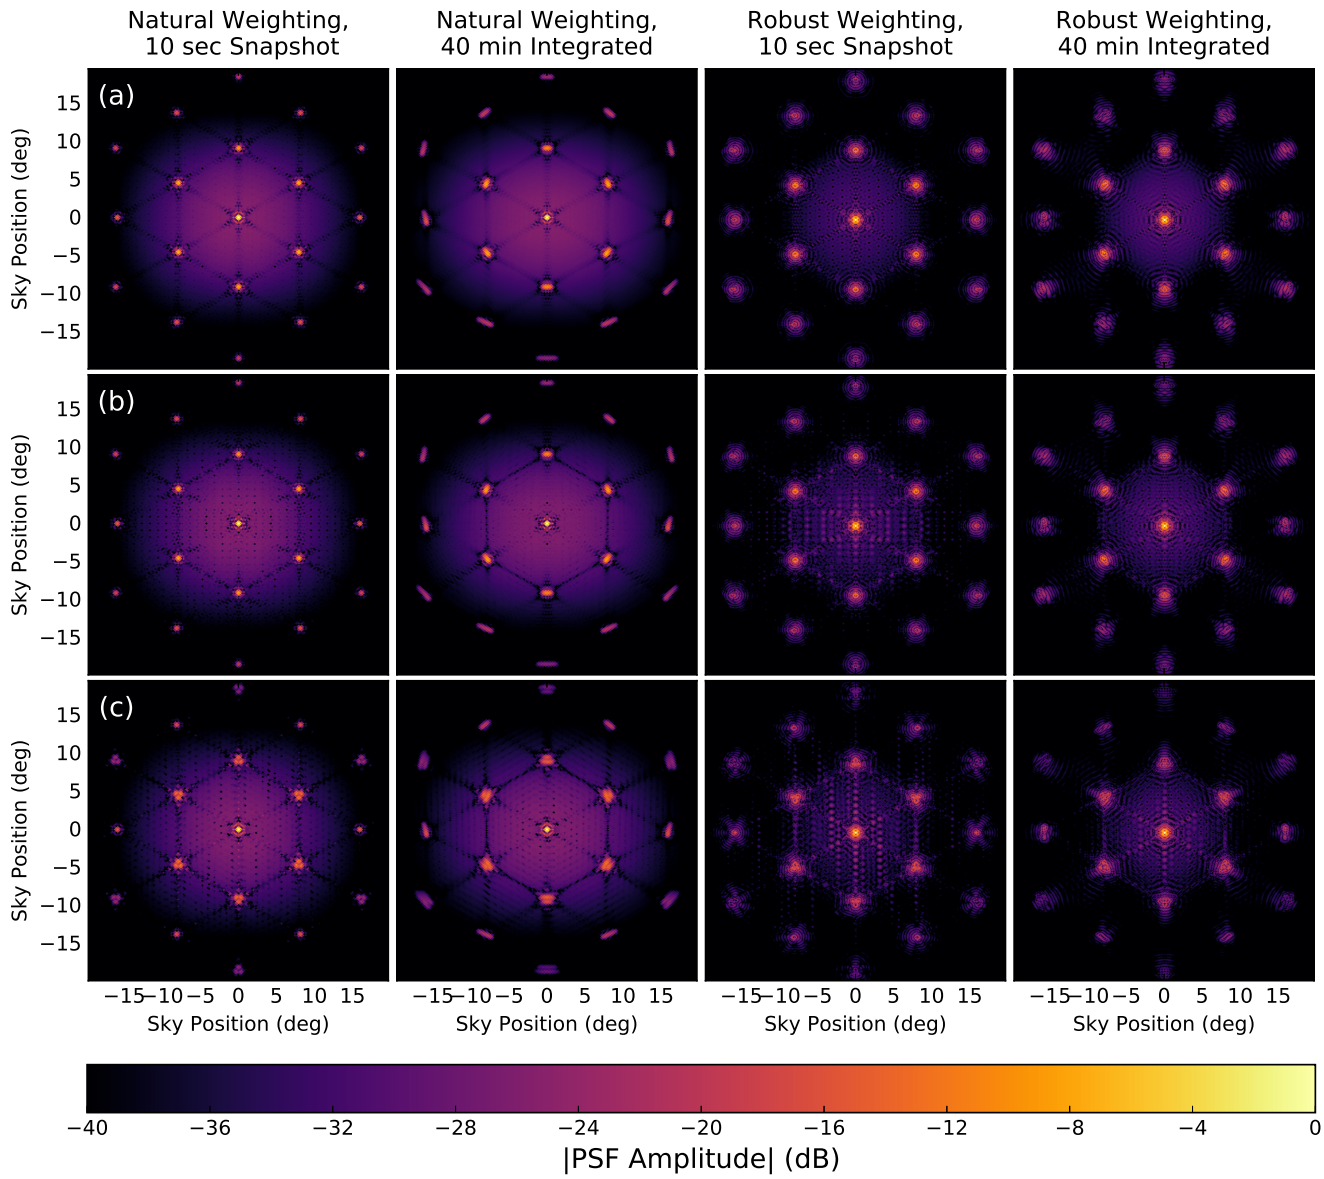

Figure 5. Point spread functions for a source at zenith, imaged at 150 MHz with all three array configurations in Figure 3. We show both single 10 second snapshots and full 40 minute integrations, corresponding to the passage of the field center through the primary beam. As we would expect, the sidelobes of the PSF or "synthesized beam" are somewhat suppressed and smeared out in configuration (c). The effect is small in the naturally weighted PSFs, since most of the sensitivity in all configurations is concentrated on the baselines in the hexagonal grid created by the core. We also include a "robust" weighting scheme, where no baseline is weighted more than 50 times more heavily than any other, in order to further highlight the sidelobe suppression that configuration (c) achieves with its relatively equitable distribution of sensitivity between on-grid and off-grid baselines.

tions (Dillon et al. 2015b).

3.3.1. *Widefield Mapmaking*

First we focus on how the core configuration (see Figure 1) affects widefield mapmaking. We therefore plot the eigenvalue spectra of all three array configurations in Figure 6. In order to overcome the numerical difficult of evaluating $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ and its eigenspectrum, we pixelize the sky adaptively, using the hierarchical property of HEALPix to develop a pixelization scheme that roughly preserves equal primary beam power per pixel. This scheme ranges from $N_{\rm side} = 64$ around zenith to $N_{\rm side} = 4$ near the horizon. We perform this analysis for both a single snapshot and 40 minutes of rotation synthesis. We then calculate the eigensystem, normalizing by the number of antennas to make the comparison fair.

At first glance, the eigenvalues—which encode the amount of information in each mode—are very similar, especially after Earth rotation synthesis. The biggest differences are revealed only in the case of instantaneous

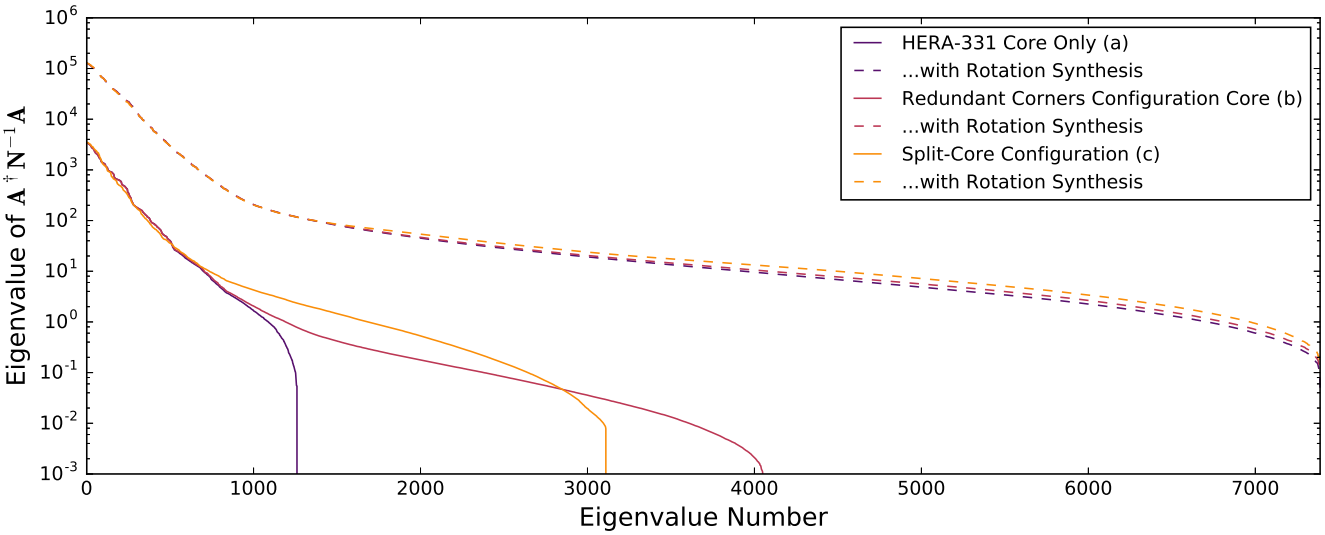

Figure 6. To assess the instantaneous and time-integrated widefield imaging capabilities of the array cores in Figure 1, we calculate the eigenvalues of $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$, the noise-weighted instrument response, for each. The amplitude of the eigenvalues of this matrix tells us how well different modes in the map are measured. It is also the matrix we would have to invert if we want to set the PSF to the identity. In this case, we examine a dynamically pixelized sky designed with HEALPix $N_{\rm side} = 64$ inside the main lobe of the primary beam and decreasing resolution at lower altitudes to maintain roughly constant beam power per pixel. Results are rescaled by the number of antennas to make the comparison fair; the units are arbitrary. We find that the addition of dithered antennas and off-grid baselines improves instantaneous mapmaking considerably, but the difference between configurations after Earth rotation synthesis (which of course adds information) is slight. Configuration (a), the unmodified hexagonal core, measures the fewest modes instantaneously. Configuration (b) can measure more independent modes than configuration (c), but with less sensitivity. This makes sense given the calculated instantaneous redundancy (Figure 2).

imaging, where we see a sharp cutoff for configuration (a) and a factor of a few difference between (b) and (c) in some of the more poorly measured eigenmodes.[6] We expect that the performance of configuration (c) is due to the more even distribution of on-grid and off-grid baselines (see Figure 2).

To verify our intuition about the meaning of these eigenvalue spectra, we should look at the eigenvectors of the single snapshot. It is difficult to interpret or compare individual eigenvectors. Instead, we want to look at each $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ matrix in two different subspaces—one corresponding to the first 630 eigenvalues where, in Figure 6, the spectra are very similar and the other corresponding to the rest where they differ the most. In Figure 7 we show the diagonals of $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ for each array configuration, projected into the two orthogonal subspaces. One can think of these matrices as $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ where either the first 630 or the rest of the eigenvalues are set to zero.

Roughly speaking, Figure 7 confirms our intuition that the highest information modes are those concentrated in the main lobe of the primary beam. That makes sense, since that is where the telescope is most sensitive. However, when we compare configuration (c) to either (a) or (b), we see a much clearer separation of information. The modes that contain information about the widefield are not the same modes as those that contain information about the primary field of view. It is not that configuration (c) is more sensitive to the modes deep in the primary beam. In fact, when when we add these two projected matrices back together, the results for all three arrays are nearly identical. Rather, what this result tells us is that we can make maps that more easily separate sources in the sidelobes from signal in the main lobe of the primary beam. As Thyagarajan et al. (2015c) and Pober et al. (2016) argue, these are the sources most likely to leak into the EoR window. The exact level of that isolation depends both on particular field of view and the observing frequency, both of which set the relative contributions of foregrounds in the main lobe and in the sidelobes of the primary beam.

3.3.2. *High-resolution Mapmaking*

We perform this analysis again on the full arrays (with outriggers) and obtain very similar results. In Figure 8 we assess the high-resolution mapmaking capability of each configuration by examining $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ for a 10° facet around at the zenith at the midpoint of the observation at high resolution ($N_{\rm side} = 256$).

[6] To be clear, each array corresponds to a different $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ matrix. Although their eigenmodes are similar, it is misleading to directly compare the Nth eigenvalue without also considering the modes they correspond to.

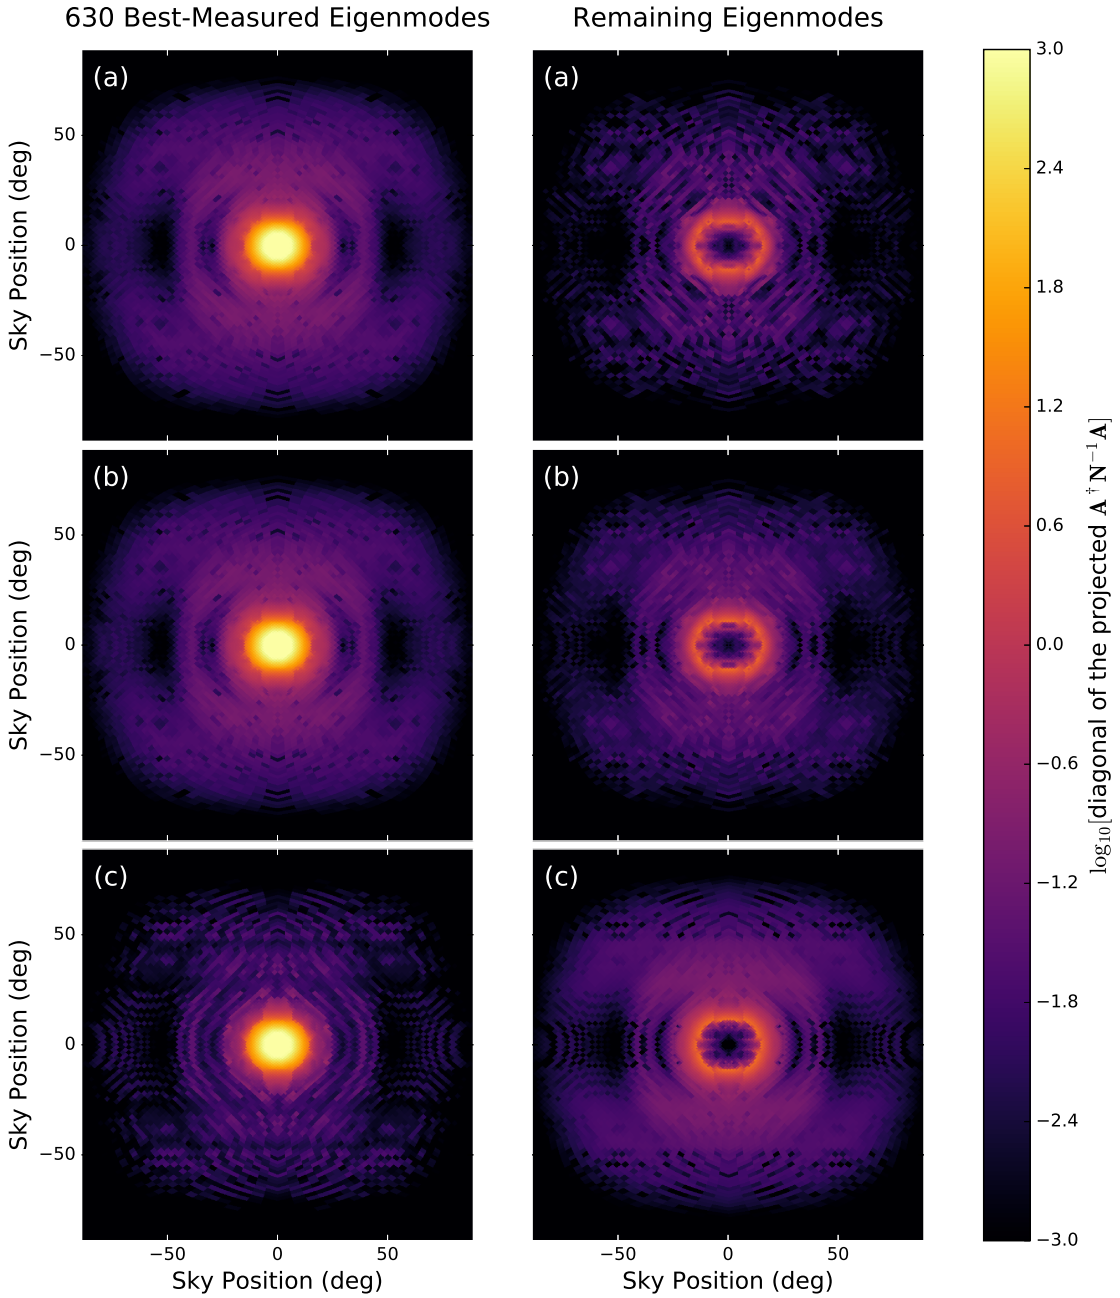

Figure 7. We can better understand the relative snapshot mapmaking performance of the array configurations in Figure 6 by looking at the diagonal elements of $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ (which is also the inverse noise covariance of the array) after it's been split into two orthogonal subspaces—one the subspace spanned by the eigenvectors corresponding the 630 largest eigenvalues, the other by the rest of the eigenvectors. 630 is where the split-core configuration (c) begins to outperform the others (see Figure 6). While the effect is modest, we see that the measurements of configuration (c) exhibit a greater separation between modes in the main lobe of the primary beam and modes in the sidelobes. This can be understood as a consequence of the brighter grating lobes of the PSF we saw in Figure 5 for configurations (a) and (b), which correlate modes inside and outside the main lobe of the primary beam.

We find that configurations (a) and (b), with their more compact cores, have more sensitivity to the best measured modes. By contrast, configuration (c) spreads its sensitivity out over a much larger number of independent modes. At very high eigenvalue number, we actually see configuration (a) performing well, due to its greater number of independently measured long baselines, although these modes are many orders of magnitude less sensitive than the best measured modes. Interestingly, even when rotation synthesis is included, the differences between the arrays are more persistent than they were in the widefield case. This is likely because the arrays differ more in their uv-coverage at long baselines than at short baselines, where the main hexagonal grid dominates.

Once again, by splitting the eigenspace of $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ for all three arrays at the eigenvalue crossing point of $N_\lambda \approx 1000$, we can get a sense of where the information comes from. Just as in Figure 7, we again see in Figure 9 that configuration (c) has a better separation of information between the center and the edge of the main lobe of the primary beam. Configuration (a), with almost no off-grid baselines, cannot separate information between zenith and the widefield very well at all.

Although the effects on mapmaking are modest, especially after Earth rotation synthesis, the separability of information between the main lobe of the primary beam and the sidelobes is interesting and potentially quite useful. In order to subtract foregrounds with sufficient fidelity to work within the wedge, we need precise subtraction of contaminants far from the main field of view (Pober et al. 2016; Thyagarajan et al. 2015c). The ability to instantaneously distinguish between foreground sidelobes and EoR signal is useful, especially when foreground positions and fluxes change due to the ionosphere (Vedantham & Koopmans 2015, 2016; Yatawatta et al. 2013; Ewall-Wice et al. 2016b). Likewise, the cleaner separation between modes will likely improve traditional sky model-based calibration or hybrid calibration schemes that rely on a sky model for bandpass calibration.

4. REDUNDANT CALIBRATABILITY

Of course, good mapmaking is only one route to good calibration. If we cannot calibrate our instrument well, we risk introducing spectral structure into our visibilities and our maps—spectral structure that, when multiplied by our bright foregrounds, may contaminate the EoR window. Of the first generation of 21 cm observatories, PAPER has so far produced the most constraining limits on the 21 cm power spectrum, despite having the least collecting area. It owes that success to its relatively simple design, its tightly focused "experimental" approach that concentrates sensitivity on just a few baselines, and its use of redundant baseline calibration.

In this section, we begin by reviewing the mathematical formalism that underlies redundant baseline calibration in Section 4.1. Then, in Section 4.2 we assess the redundant calibratability of the array configurations proposed and examined above. In particular, we show how even the off-grid antennas and outriggers can be redundantly calibrated.

4.1. *Redundant Calibration Review*

Fundamentally, the calibration problem for a radio interferometer can be expressed as[7]

$$V_{ij}^{\text{measured}} = g_i g_j^* V_{i-j}^{\text{true}} + n_{ij}. \tag{8}$$

The visibility we measure, V_{ij}^{measured}, is related to the true visibility by the product of the complex gain g_i of each antenna, plus some noise. We write the true visibility in shorthand as V_{i-j}^{true} to express the idea that all that matters is the baseline, the separation between antennas i and j, rather than the specific antennas involved.

The key idea behind redundant baseline calibration is that the number of numbers we want to estimate, the gains and true visibilities, is much smaller than the number of measurements we actually make. We have an overdetermined system. If we can linearize Equation 8, then we can write down a vector expression in the form of Equation 2 that relates all the visibilities measured to all the gains and true visibilities. Then we can use exactly the same statistical machinery to solve for both the gains and visibilities simultaneously.

Liu et al. (2010) propose two ways of doing that. The simpler way, called `logcal`, linearizes Equation 8 by taking the natural logarithm. If we express the complex gains as $g_j \equiv \exp[\eta_j + i\phi_j]$, then it follows that

$$\ln\left(V_{ij}^{\text{measured}}\right) = \eta_i + \eta_j + i\phi_i - i\phi_j + \ln\left(V_{i-j}^{\text{true}}\right) + n'_{ij}. \tag{9}$$

The precise form and statistics of n'_{ij} are discussed extensively in Liu et al. (2010). This can be rewritten in the same form as Equation 2 by treating $\mathbf{y}$ as the set of measurements, $\mathbf{x}$ as a vector that contains both the gains and the true visibilities, and $\mathbf{A}$ as the matrix whose rows specify which pairs of antennas with which gains measure which visibility. In this form the real and imaginary parts of the gains and visibilities, produce two

[7] This leaves aside the question of "direction-dependent calibration," which is essentially primary beam calibration. In this work, we assume our primary beams are the same. While it is possible to take antenna-dependent beams into account when making maps (Dillon et al. 2015b), it is not clear how much antenna-to-antenna variation (Neben et al. 2015, 2016a) affects redundant baseline calibration. It should also be pointed out that the sub-grid—and therefore sub-aperture—sampling enabled by all our configurations makes at least some kind of direction-dependent calibration possible using redundant information. Both questions are beyond the scope of this paper and are left for future work.

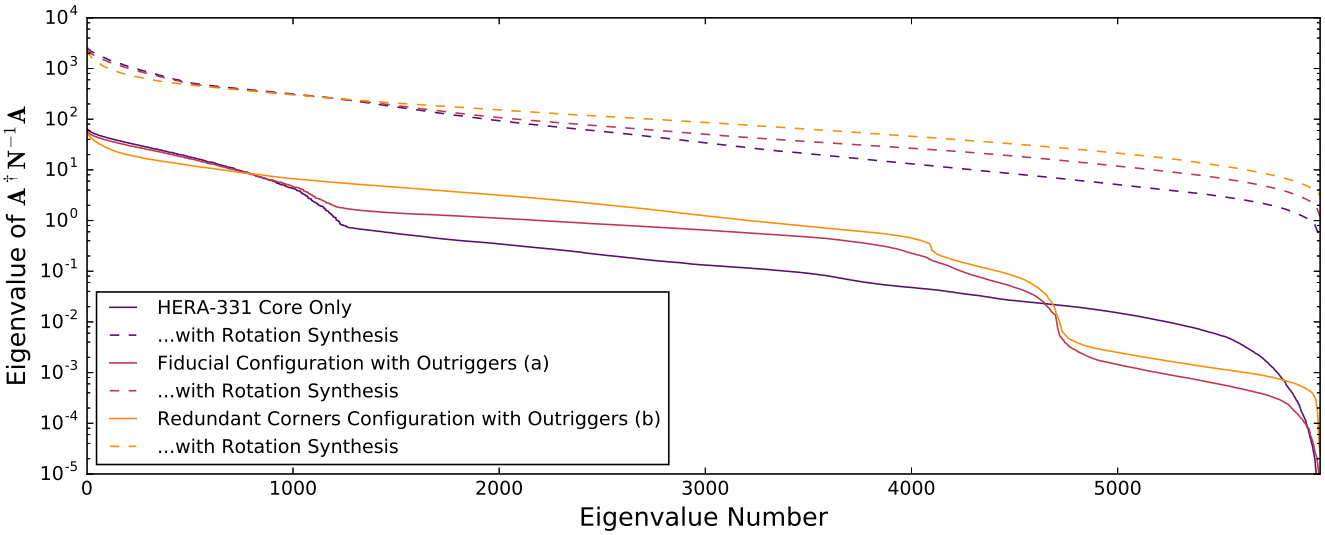

Figure 8. Here we examine the instantaneous and time-integrated high-resolution imaging capabilities of the arrays (with outriggers) shown in Figure 3. Unlike in Figure 6, we calculate the PSF for only a 10° facet, ignoring correlations with the rest of the sky, at HEALPix $N_{\rm side} = 256$. We find that (b) and (c) concentrate their sensitivity on a relatively small number of modes while (c) spreads out the sensitivity. This effect is especially evident for instantaneous imaging, though still pretty modest. Since $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ is proportional to observing time, any shortcomings here can be erased by doubling the length of the observational campaign.

different sets of equations.

It turns out that `logcal`'s solutions are slightly biased. To rectify that, Liu et al. (2010) also developed `lincal`, a method of linearizing Equation 8 by Taylor expanding it and then solving for gains and visibilities iteratively. Both methods were implemented for Zheng et al. (2014) and were integrated into the PAPER analysis pipeline (Ali et al. 2015), with `logcal` providing `lincal` a good starting guess.[8]

In our case, we are interested in how the array configuration affects the errors on redundant calibration. For simplicity, we look at the expected gain errors predicted by `logcal`. If we form an unbiased estimator of the gains and true visibilities, then the gain/visibility covariance, $\mathbf{\Sigma}$, is given by

$$\mathbf{\Sigma} = \left[\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}\right]^{-1}. \tag{10}$$

In the case where all baselines measure visibilities with the same magnitude[9] and all antennas have the same noise level, Liu et al. (2010) show that $\mathbf{\Sigma}$ is reduced to

$$\mathbf{\Sigma} \longrightarrow \frac{1}{(\mathrm{S/N})^2}\left[\mathbf{A}^\dagger\mathbf{A}\right]^{-1}, \tag{11}$$

where S/N is the S/N of the sky signal (i.e., foregrounds), not just the 21 cm signal. This leads us to a new metric for assessing array configurations by looking at the gain errors predicted by the diagonal of $\left[\mathbf{A}^\dagger\mathbf{A}\right]^{-1}$ It can be shown that the $\Delta\eta$s we predict are simply the gain errors up to third order in $\Delta\eta$:

$$\Delta|g_i| = \Delta\eta_i + \mathcal{O}(\Delta\eta^3). \tag{12}$$

There remains one important complication related to the essential nature of what redundant calibration can and cannot do. As it turns out, $\mathbf{A}^\dagger\mathbf{A}$ is actually never invertible. There must always be at least one zero-eigenvalue because there is always an overall amplitude degeneracy in this system. We can increase the gains by some factor, decrease the true visibilities by that factor squared, and still get the same measurements. For the imaginary half of `logcal`, it turns out that there are always at least three zero-eigenvalues, one overall phase and two tip/tilt terms related to the orientation of the array.[10] This limitation is discussed extensively in Zheng et al. (2014).

This is not to say that redundant baseline calibration is not helpful. It reduces the problem of calibrating a gain and a phase for every antenna at every frequency to calibrating just a few numbers per frequency for the entire array. The latter problem is less noisy, less sensi-

[8] Both methods have been integrated into a single package and are available at `https://github.com/jeffzhen/omnical`

[9] This would be the case for a single dominant point source. At the frequencies and fields of view relevant to 21 cm cosmology, that approximation would generally be a bad one. However, it is an illustrative limit to consider in our case when most visibilities are still the same order of magnitude.

[10] In `lincal` or any other approach to redundant calibration, there will always be at least four zero-eigenvalues due to these degeneracies.

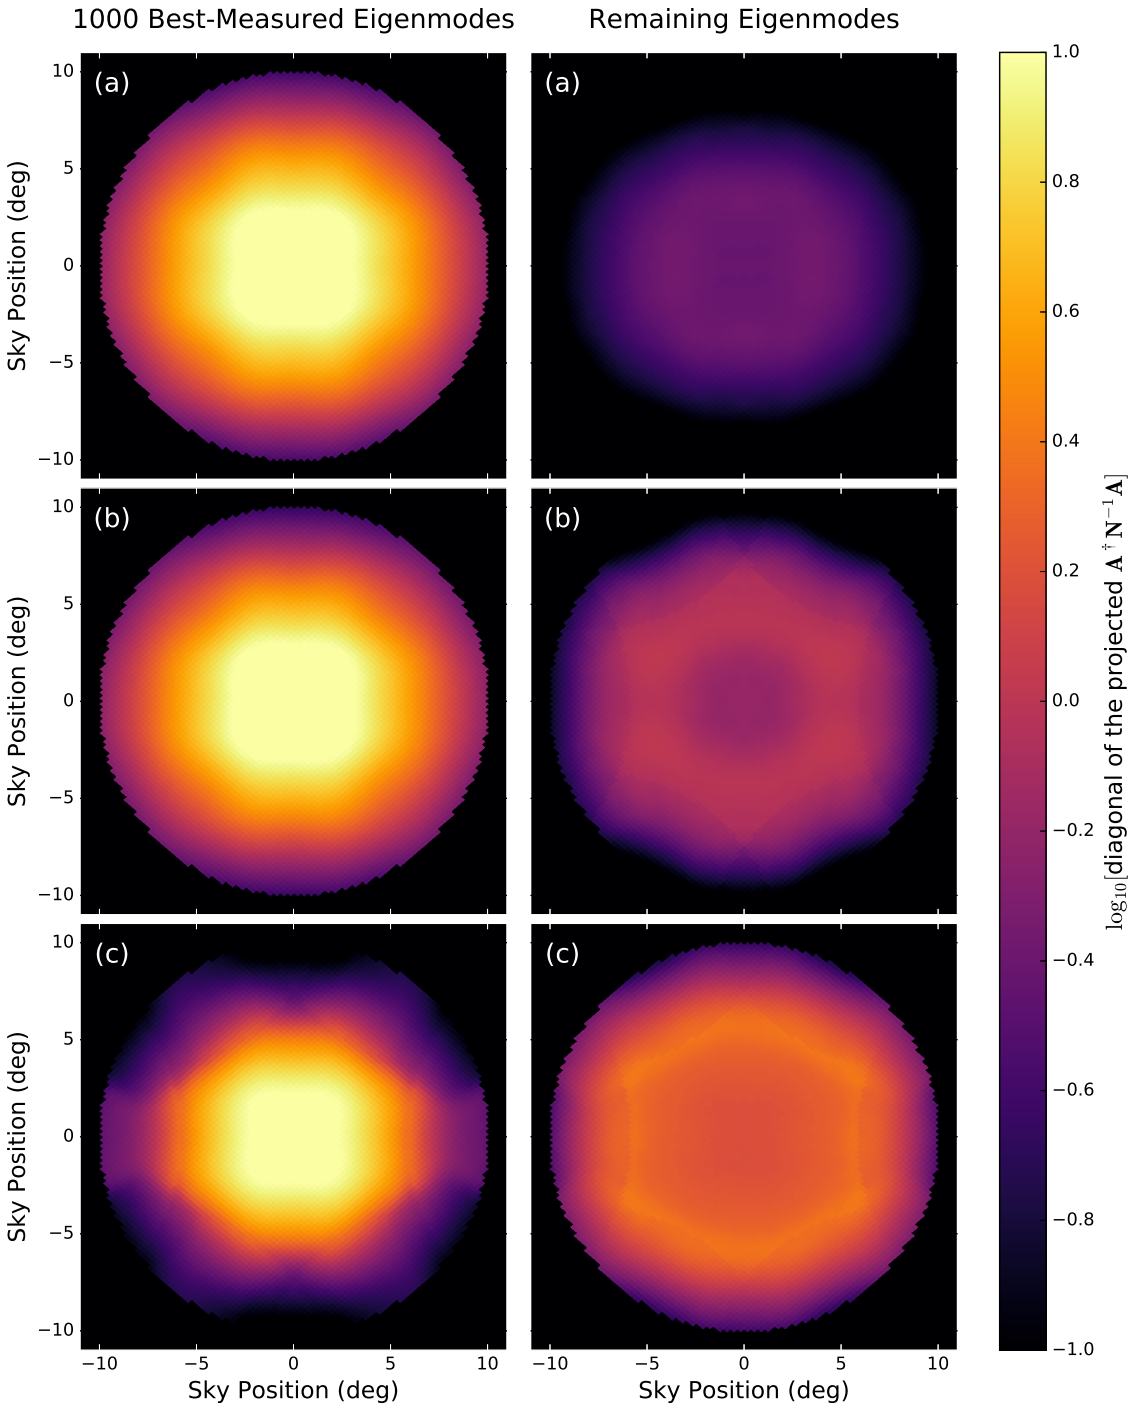

Figure 9. Just as in Figure 7, we can understand the significance of the main lobe, high-resolution snapshot eigenvalue spectrum of $\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}$ in Figure 8 by looking that the diagonal elements of the matrix when it has been separated into two orthogonal subspaces. That separation—between the first 1000 eigenvalues and the rest—is again set by the crossing point where array configuration (c) begins out outperform the others. Once again, we see a cleaner separation between information in the center of the field and in the wide field. This supports the idea that configuration (c) sacrifices some sensitivity to the main lobe in favor of increased information about the sidelobes, which makes intuitive sense since the split-core configuration has observations more evenly distributed between on-grid and off-grid baselines (see Figures 2 and 4).

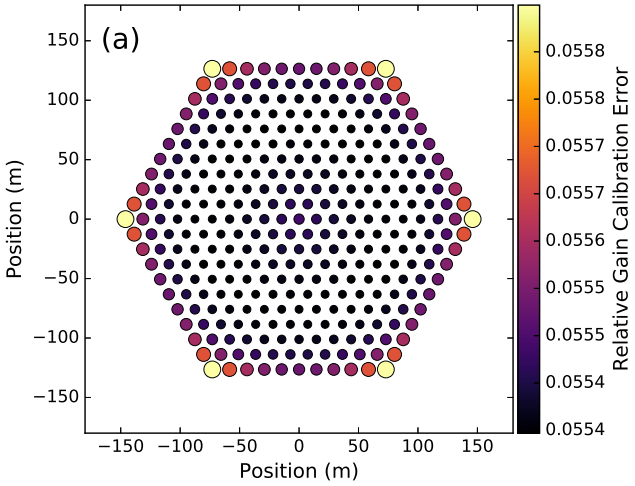

Figure 10. Here we show the expected error on the calibration (using both color and relative marker sizes) of antenna gains, relative to the S/N, for the core of configuration (a), the 331-element HERA hexagon. Using redundant baseline calibration, very precise gain solutions are attainable throughout. This calculation assumes that all baselines measure visibilities with the same amplitude, although in practice, amplitudes vary with baseline length, baseline orientation, and time. As in Liu et al. (2010), we note that edge antennas are generally slightly less calibratable than interior antennas. Interestingly, the most calibratable antennas are off-center; this is because those antennas participate in a greater number of unique baseline types.

tive to the accuracy of the sky model, and therefore far less error-prone. Fortunately, the gain errors we examine are insensitive to the overall amplitude, since they are normalized by the S/N. We merely need to take the pseudoinverse of $\mathbf{A}^\dagger\mathbf{A}$ to obtain $\mathbf{\Sigma}$ (Dillon et al. 2013).

4.2. *Redundant Calibratability Results*

Before we look at the three array configurations considered above, let us try to develop some intuition for intuition for this formalism. In Figure 10, we plot the gain errors expected from redundant calibration of a 331-element solid hexagonal HERA core, the core of configuration (a). Plotted are the gain errors we would expect if each visibility had the same amplitude and were measured with a foreground S/N of 1. The errors depend inverse-linearly on the S/N.

For a highly redundant configuration like the solid HERA core, the antenna-to-antenna variation in gain errors is very small. In these cases, the relative gain errors are all roughly just $(N_{\rm antennas})^{-1/2}$ (Liu et al. 2010). However, we can see some illustrative structure in Figure 10. The edge and especially corner antennas are are slightly more difficult to calibrate—they are involved in fewer highly redundant baselines than interior antennas. Likewise, antennas in the very center are involved in fewer kinds of baselines and thus also see slightly higher error. Both the redundancy of the baselines an antenna is involved in and the centrality of those baselines to the larger connected graph of redundant baselines affect the calibratability of each antenna. Regardless, for a 331-element core these variations are small in comparison to what we see for more complex array designs; the key comparison we wish to make is how well we can redundantly calibrate our full arrays relative to the a solid core.

After calculating the `logcal` $\mathbf{A}^\dagger\mathbf{A}$ for each of the three arrays we have considered above, configuration (a) sticks out right away. The imaginary version $\mathbf{A}^\dagger\mathbf{A}$ should have three zero-eigenvalue; it has five. Two extra phase factors are not redundantly calibratable because the outrigger baselines do not overlap as they do in configuration (b). While the core of configuration (a) can be redundantly calibrated, its outriggers cannot. We therefore cannot predict the errors from redundant calibration for this array.

Configurations (b) and (c), however, were both designed to be redundantly calibratable.[11] They have the correct number of zero-eigenvalues, meaning that taking the pseudoinverse of $\mathbf{A}^\dagger\mathbf{A}$ is appropriate. In Figure 11 we show the expected errors on antenna gains when calibrating the cores redundantly and when calibrating the whole array redundantly.

When calibrating just the cores, the addition of the 12 off-grid corner antennas in configuration (b) leads to higher gain errors not just in the off-grid antennas but also, surprisingly, in the hexagonal core by about 20%. By contrast, configuration (c) sees raised errors generally by less than 1% compared to a solid HERA core. This is because the core of configuration (b) probes many more baselines than a simple HERA core, but most of them are sparsely sampled. A large number of baselines are only weakly connected to the rest, leading to higher errors throughout. Of course, if one wishes to give up on redundant calibration of the off-grid antennas, one can achieve the same errors as in Figure 10. However, then one must resort to another strategy (i.e. one that requires a sky model) to calibrate the remaining antennas (which affect all the off-grid visibilities).

Those effects are mitigated somewhat by the addition of redundant outriggers that decrease the discrepancy between the array cores to a 5% difference in calibration errors. On the other hand, the most distant outriggers in configuration (b) have approximately 50% larger gain errors than the most furthest-flung outriggers in configuration (c). Simply put, the evenness of the distribution of redundancy across the baselines probed, and thus the connectedness of antennas through that network of re-

[11] Even in the event that a single antenna breaks, they can still be redundantly calibrated.

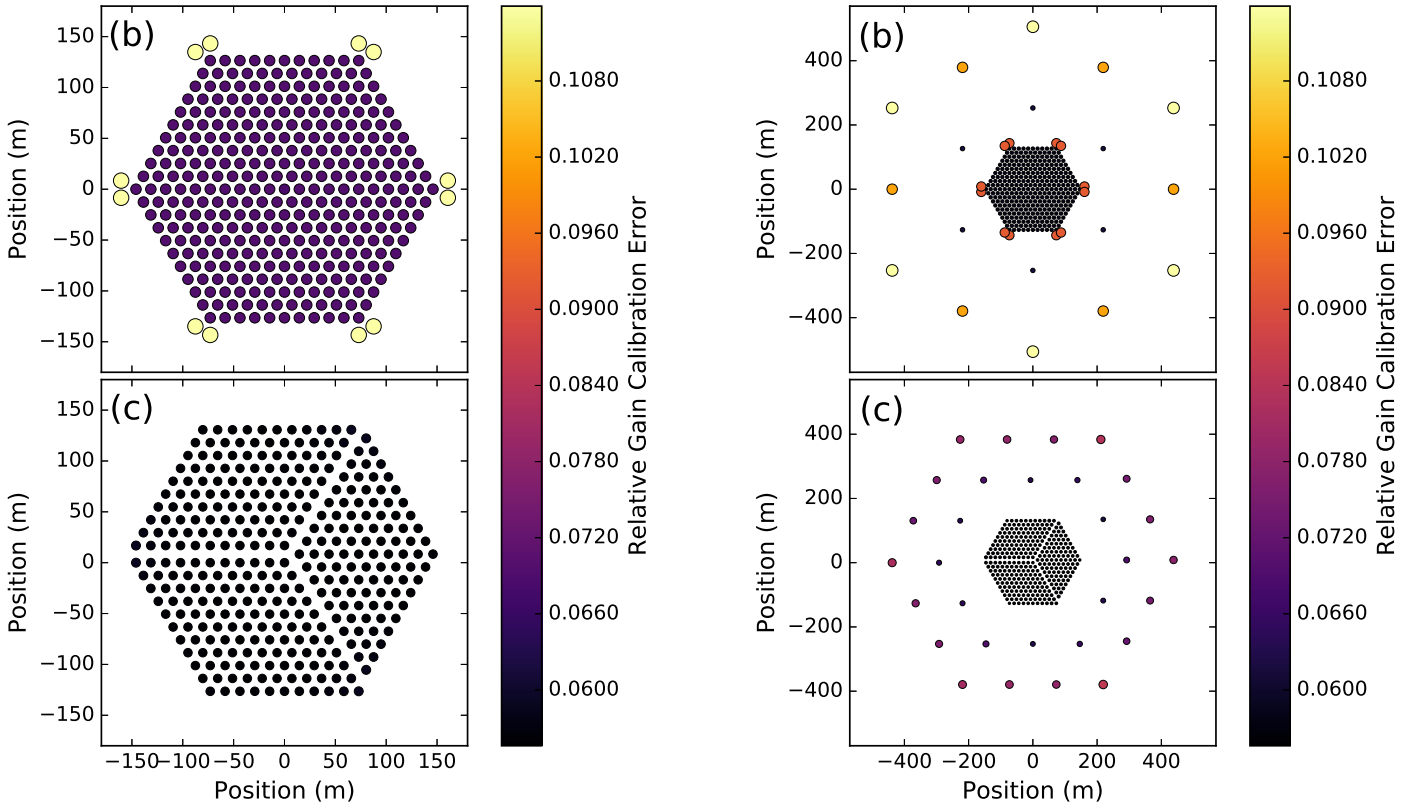

Figure 11. Expected gain errors relative to S/N (shown with both color and relative marker size) for configurations (b) and (c). On the left we show the relative gain errors of the cores alone, on the right we show all the errors when calibrating the full arrays using redundant calibration. When only calibrating the cores, configuration (b) shows increased gain errors on the whole array by about 20%. The dithered corner antennas, in particular, have roughly double the error of the rest of the array, meaning that all off-grid baselines will be effectively noisier by a factor of 2 (in the visibilities). Configuration (c) has almost the same errors in the core as an unmodified hexagonal core. When outriggers are added, the differences between the cores are reduced to only about 5%. However, the furthest outriggers have considerably larger error in configuration (b) than configuration (c). Note that configuration (a) is excluded here because it is not redundantly calibratable with its outriggers.

dundant baselines, makes configuration (c) more redundantly calibratable.

5. SUMMARY OF ARRAY DESIGN LESSONS

In this work, we examined how minor modifications to interferometric array configurations with dense cores optimized for 21 cm cosmology can both produce useful images and be redundantly calibrated, even though they are designed for foreground avoidance and maximal sensitivity in the EoR window. Based on Dillon et al. (2015b) and Liu et al. (2010), we quantified these effects by examining the information content or, relatedly, the expected error. These are two new ways to think about array configurations and we expect them to be broadly useful for designing future 21 cm arrays that rely on dense cores and redundant calibration.

The differences between the arrays we examined for mapmaking capability and redundant calibratability were relatively small, although we can attribute this to the restriction that we maintain a high-sensitivity, dense core optimized for foreground avoidance—currently the most robust and promising approach for accessing the high-redshift 21 cm signal. Despite that, we can extrapolate a number of relevant lessons for future telescopes:

1. As expected, adding "off-grid" baselines improves instantaneous widefield imaging and the separation between modes in the main lobe of the primary beam and those further from the zenith.

2. A redundantly calibratable array need not necessarily be bad for imaging. A redundantly calibratable array can still have dithered, off-grid baselines. It can still have distant outriggers. And even though redundant calibration relies on measuring only a relatively modest number of unique visibilities at any given time, we can get back sensitivity to a large number of distinct modes with Earth rotation synthesis.

3. Redundant calibratability does not simply depend on the number of antennas and the number of unique visibilities. The "connectedness" of the

relationship between gains and visibilities—which well-measured visibilities are most useful for constraining which gains and vice versa—is important as well. This is a reason to prefer configurations like (c) over simply adding a few off-grid antennas.

4. The linear estimator formalism that underlies both calculations in this paper is a powerful tool for assessing the expected performance of an instrument during calibration and data reduction.

In time, it would be useful to examine how these effects propagate to the scientific results we expect from instruments like HERA: 21 cm power spectra and constraints on reionization. While Greig et al. (2015) was able to perform this analysis for widely differing SKA configurations, they relied on approximate sensitivity calculations that ignore correlations between uv-cells—precisely the effect that matters when we start examining off-grid and therefore sub-aperture baselines. To our knowledge, no power spectrum analysis has been demonstrated that takes into account full noise correlations and the frequency dependence of both the primary beam and the PSF. A rigorous and more quantitative examination of array configurations and their effect on astrophysical and cosmological constraints is left to future work. However, our new techniques show how to make small sacrifices in sensitivity in order to obtain sizable improvements in the control of systematics. With better instrument simulation and power spectrum estimation, one can assess the metrics we have developed here to understand their relative effects on 21 cm cosmology.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the contributions of a number of colleagues and collaborators. Adrian Liu was extremely helpful in giving feedback and general advice about the direction of this work. Ue-Li Pen gave us ideas that eventually led to configuration (c). Jeff Zheng helped us use his `Omnical` package and double-checked the redundant calibratability of configurations (b) and (c). Likewise, Jonathan Pober helped us apply his `21cmSense` package. Miguel Morales and Bryna Hazelton gave us useful feedback on how to examine PSFs and sidelobe suppression. David DeBoer gave us feedback on the feasibility of these array configuration modifications (and more radical ones not considered in this work) and, with James Aguirre, supplied us with HERA primary beam simulations. Finally, we would like to thank Max Tegmark for originally sparking an interest in the question of array configuration optimization, starting with Tegmark & Zaldarriaga (2010), and continuing with discussions for years afterwards.

This research was conducted as part of the University of California Cosmic Dawn Initiative. The authors acknowledge support from the University of California Office of the President Multicampus Research Programs and Initiatives through award MR-15- 328388. The authors also acknowledge support form the National Science Foundation through CAREER award No. 1352519, AST grant No. 1129258, and AST grant No. 1440343.

REFERENCES

Ali, Z. S., Parsons, A. R., Zheng, H., et al. 2015, ApJ, 809, 61
Bernardi, G., Mitchell, D. A., Ord, S. M., et al. 2011, MNRAS, 413, 411
Bonaldi, A., & Brown, M. L. 2015, MNRAS, 447, 1973
Bowman, J. D., Cairns, I., Kaplan, D. L., et al. 2013, PASA, 30, 31
Braun, R. 2013, A&A, 551, A91
Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, MNRAS, 429, 165
Clesse, S., Lopez-Honorez, L., Ringeval, C., Tashiro, H., & Tytgat, M. H. G. 2012, PhRvD, 86, 123506
Datta, A., Bowman, J. D., & Carilli, C. L. 2010, ApJ, 724, 526
de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, MNRAS, 388, 247
DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2016, ArXiv e-prints, arXiv:1606.07473
Dillon, J. S., Liu, A., & Tegmark, M. 2013, PhRvD, 87, 043005
Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, PhRvD, 89, 023002
Dillon, J. S., Neben, A. R., Hewitt, J. N., et al. 2015a, PhRvD, 91, 123011
Dillon, J. S., Tegmark, M., Liu, A., et al. 2015b, PhRvD, 91, 023002
Ewall-Wice, A., Hewitt, J., Mesinger, A., et al. 2016a, MNRAS, 458, 2710
Ewall-Wice, A., Dillon, J. S., Hewitt, J. N., et al. 2016b, MNRAS, arXiv:1605.00016
Ewall-Wice, A., Bradley, R., DeBoer, D., et al. 2016c, ApJ, submitted, arXiv:1602.06277
Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, PhR, 433, 181
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Greig, B., & Mesinger, A. 2015, MNRAS, 449, 4246
Greig, B., Mesinger, A., & Koopmans, L. V. E. 2015, MNRAS, submitted, arXiv:1509.03312
Hazelton, B. J., Morales, M. F., & Sullivan, I. S. 2013, ApJ, 770, 156
Jacobs, D. C., Pober, J. C., Parsons, A. R., et al. 2015, ApJ, 801, 51
Koopmans, L., Pritchard, J., Mellema, G., et al. 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), 1
Liu, A., Parsons, A. R., & Trott, C. M. 2014a, PhRvD, 90, 023018
—. 2014b, PhRvD, 90, 023019
Liu, A., Tegmark, M., Morrison, S., Lutomirski, A., & Zaldarriaga, M. 2010, MNRAS, 408, 1029
Loeb, A., & Furlanetto, S. R. 2013, The First Galaxies In The Universe (Princeton University Press, Princeton, NJ)
Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. 2009, IEEE Proceedings, 97, 1497

Mao, Y., Tegmark, M., McQuinn, M., Zaldarriaga, M., & Zahn, O. 2008, PhRvD, 78, 023529
McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, ApJ, 653, 815
Mesinger, A., Furlanetto, S., & Cen, R. 2011, MNRAS, 411, 955
Morales, M. F. 2011, PASP, 123, 1265
Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, ApJ, 752, 137
Morales, M. F., & Wyithe, J. S. B. 2010, ARA&A, 48, 127
Neben, A. R., Bradley, R. F., Hewitt, J. N., et al. 2015, Radio Science, 50, 614
Neben, A. R., Hewitt, J. N., Bradley, R. F., et al. 2016a, ApJ, 820, 44
Neben, A. R., Bradley, R. F., Hewitt, J. N., et al. 2016b, ApJ, submitted, arXiv:1602.03887
Paciga, G., Albert, J. G., Bandura, K., et al. 2013, MNRAS, 433, 639
Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012a, ApJ, 753, 81
Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012b, ApJ, 756, 165
Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, AJ, 139, 1468
Parsons, A. R., Liu, A., Aguirre, J. E., et al. 2014, ApJ, 788, 106
Pen, U.-L., Wu, X.-P., & Peterson, J. 2004, ArXiv Astrophysics e-prints, astro-ph/0404083
Pober, J. 2015, Updated 21cm Experiment Sensitivities, Tech. Rep. 4, HERA, http://reionization.org/science/memos
Pober, J. C., Parsons, A. R., Aguirre, J. E., et al. 2013, ApJL, 768, L36
Pober, J. C., Liu, A., Dillon, J. S., et al. 2014, ApJ, 782, 66
Pober, J. C., Hazelton, B. J., Beardsley, A. P., et al. 2016, ApJ, 819, 8
Pritchard, J. R., & Loeb, A. 2012, Reports on Progress in Physics, 75, 086901
Ram Marthi, V., & Chengalur, J. 2014, MNRAS, 437, 524
Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2015, PhRvD, 91, 083514
Sullivan, I. S., Morales, M. F., Hazelton, B. J., et al. 2012, ApJ, 759, 17
Tegmark, M. 1997, ApJL, 480, L87
Tegmark, M., & Zaldarriaga, M. 2009, PhRvD, 79, 083530
—. 2010, PhRvD, 82, 103501
Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, Interferometry and Synthesis in Radio Astronomy; 2nd ed. (Weinheim: Wiley-VCH)
Thyagarajan, N., Beardsley, A. P., Bowman, J. D., & Morales, M. F. 2015a, MNRAS, submitted, arXiv:1510.08318
Thyagarajan, N., Parsons, A., DeBoer, D., et al. 2016, ApJ, submitted, arXiv:1603.08958
Thyagarajan, N., Udaya Shankar, N., Subrahmanyan, R., et al. 2013, ApJ, 776, 6
Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015b, ApJL, 807, L28
—. 2015c, ApJ, 804, 14
Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, PASA, 30, 7
Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, ApJ, 757, 101
Trott, C. M., Pindor, B., Procopio, P., et al. 2016, ApJ, 818, 139
van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2
Vedantham, H., Udaya Shankar, N., & Subrahmanyan, R. 2012, ApJ, 745, 176
Vedantham, H. K., & Koopmans, L. V. E. 2015, MNRAS, 453, 925
—. 2016, MNRAS, 458, 3099
Wieringa, M. H. 1992, Experimental Astronomy, 2, 203
Wijnholds, S. J., Bregman, J. D., & van Ardenne, A. 2011, Radio Science, 46, RS0F07
Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, A&A, 550, A136
Zaroubi, S. 2013, in Astrophysics and Space Science Library, Vol. 396, The First Galaxies, ed. T. Wiklind, B. Mobasher, & V. Bromm, 45
Zheng, H., Tegmark, M., Buza, V., et al. 2014, MNRAS, 445, 1084